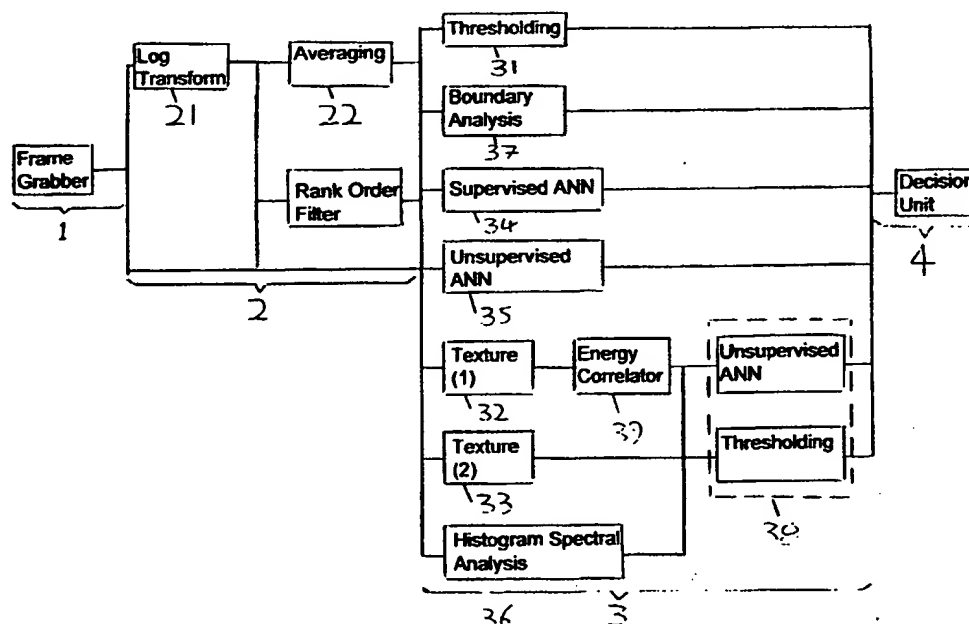




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06T 7/00		A1	(11) International Publication Number: WO 95/16247
			(43) International Publication Date: 15 June 1995 (15.06.95)
(21) International Application Number: PCT/GB94/02698 (22) International Filing Date: 9 December 1994 (09.12.94) (30) Priority Data: 9325243.5 9 December 1993 (09.12.93) GB (71) Applicant (for all designated States except US): THE ROYAL HOLLOWAY AND BEDFORD NEW COLLEGE [GB/GB]; Egham, Surrey TW20 0EX (GB). (72) Inventors; and (75) Inventors/Applicants (for US only): HANNAH, Ian [GB/GB]; 47 Larksfield, Englefield Green, Surrey TW20 0RD (GB). DAVIES, Emllyn, Roy [GB/GB]; 3 Barn Way, Englefield Green, Surrey TW20 0QU (GB). PATEL, Devesh [IN/GB]; 9 Rosebury Vale, Ruislip, Middlesex HA4 6AQ (GB). (74) Agents: CRAWFORD, Andrew, Birkby et al.; A.A. Thornton & Co., Northumberland House, 303-306 High Holborn, London WC1V 7LE (GB).			(81) Designated States: GB, JP, US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: AUTOMATIC MONITORING SYSTEM



(57) Abstract

An automatic monitoring system for detecting foreign objects or defects affecting a substrate, the system comprising imaging means for obtaining an image of the substrate, image processing means comprising a plurality of parallel image processing paths applying different analytical techniques to the image, and decision means receiving inputs from the processing paths which decides whether or not the imaged substrate is affected by a foreign object or defect. In a preferred embodiment of the invention, the monitoring system employs trainable methods involving the use of artificial neural networks either supervised or unsupervised and cognate statistical pattern recognition methods.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

- 1 -

AUTOMATIC MONITORING SYSTEM

The present invention relates to an automatic monitoring system and, more particularly, to such a system adapted to detect foreign objects or defects affecting a substrate.

The system of the present invention has been developed in the context of monitoring bags of foodstuffs for the presence of foreign objects. However, it is to be understood that embodiments of the invention are more widely applicable to monitoring food products (including liquids) whether in bags (or other containers) or not, e.g. they may be in free flow on a conveyor or in fluidised flow in a pipe (as especially in the case of grain or flour). Similarly, the monitored products may be other than food products. For example, they may be agricultural (detection of contamination in harvested cereals by both internal and external infestation and extraneous matter such as bird droppings, rodent pellets and fungi including ergot) or raw materials of various sorts (e.g. lumps of coal carried on a conveyor), or textiles, paper, liquids etc. Furthermore in addition to enabling detection of foreign bodies, the techniques of the present invention enable defects in a monitored substrate to be detected. The term 'defects' is to be taken to include damage (e.g. as it applies to the surfaces of computer discs), bruising (e.g. as it applies to commodities such as apples), or breakage, chipping, cracking, fractures, and so on (e.g. as applied to steel components or tools or welded parts). Finally, it is not deemed necessary for recognition of the exact foreign object or defect to occur in order for valid rejection to be instituted: it is commonly sufficient for the unusual to be noticed and acted upon.

The simplest systems for detecting foreign objects and defects in substrates involve human inspection. Such system are not suitable for inspection of packaged food and, in any event, can be unreliable due to loss of concentration by the monitoring personnel.

It is known to detect foreign objects in bags of foodstuffs on a production line by x-ray imaging the bags and processing the resultant grey-level images. This processing involves comparing the grey-levels of the pixels in the image with a threshold. Pixels having a grey-level above or below threshold can be considered to represent part of a foreign object. Bags of foodstuffs deemed to contain foreign objects are then rejected from the production line.

The prior art technique is of little use for detecting foreign objects made of relatively x-ray soft material ("soft contaminants") such as wood, rubber and plastics material. Furthermore, even though hard contaminants such as pieces of metal and glass can be detected by the known technique, in practice the texture of the food substrate makes the image intensities so random that small pieces of metal are hard to detect reliably. Indeed, many hard contaminants, such as pieces of glass or stone, cannot be detected by this method.

Embodiments of the present invention provide advantages over standard means of visual inspection in that they achieve a greater capability of reliably locating foreign objects and defects (FODs) than the standard intensity thresholding technique. Furthermore, embodiments of the present invention are capable of detecting soft contaminants close to boundaries in the substrate (e.g. close to the edge of a bag of vegetables).

Not only do preferred embodiments of the invention work more locally than the prior art technique by providing thresholding that adapts to the local

- 3 -

conditions existing in parts of the input images, and thereby locating the FODs without also locating excessive numbers of false alarms which would make product rejection uneconomic, but also they are capable of analysing the textures existing in many food-based (and other) images, and locating FODs in these cases too. The present invention provides an automatic monitoring system for detecting foreign objects or defects affecting a substrate, the system comprising:

imaging means for obtaining an image of at least a portion of the substrate;

image processing means comprising a plurality of parallel image processing paths applying different analytical techniques to evaluate which regions of the image corresponds to the normal substrate and which regions of the image correspond to foreign objects or defects; and

decision means receiving inputs from the parallel image processing paths and adapted to decide whether or not the imaged portion of the substrate is affected by a foreign object or defect.

The present invention further provides a method for automatically monitoring a substrate for foreign objects or defects affecting the substrate, the method comprising the steps of;

imaging at least a portion of the substrate;

processing the captured image using a plurality of parallel processing paths applying different analytical techniques to evaluate which regions of the image correspond to the normal substrate and which regions of the image correspond to foreign objects or defects; and

making a judgement as to whether or not the imaged portion of the substrate is affected by a foreign object or defect, the judging step being influenced by outputs from the parallel processing paths.

- 4 -

The present invention is expected to find most immediate application in factories and the like. In factory applications the decision means can be used to judge whether or not an imaged product should be rejected from the production line. Alternatively, a product which is deemed by the decision means to be affected by an FOD may be set aside for another purpose, for example it may be recycled, used for animal feed, labelled as satisfactory for different types of customer, etc. (referred to below as "classification").

In preferred embodiments of the invention, automatic monitoring system employs methods which are trainable, involving use of artificial neural networks (ANNs) and cognate statistical pattern recognition methods, though once the system has learned the type of product and FOD which might occur, the trainable aspect can progressively be removed. On the other hand, to retain adaptability, the trainable aspect would normally be retained to some degree in the final factory monitoring system.

The ANNs used in embodiments of the system are of three main types; supervised (capable of learning to distinguish individual FODs from the background or substrate medium); unsupervised (capable of learning the substrate medium (BSM) well enough to recognise when a FO or defect is present by the fact that the BSM parameters have locally varied from the norm); and unsupervised but providing a suitable approximation to the results of principal components analysis (PCA), so that fast adaption to new local conditions of the BSM results.

Preferred embodiments of the invention include among the parallel image processing paths means for examining image portions relating to the boundaries of packets to determine whether any FODs are in the vicinity of these boundaries. This is necessary because

- 5 -

certain of the texture-based methods are not so effective near boundaries because of the inherent region averaging processes they include. It has also been determined that incorporating intensity histogram analysis will automatically include a measure of boundary inspection, which may save implementation cost in certain practical versions of the system. This aspect of the invention makes a significant contribution to rendering embodiments of the present invention viable as apparatus/methods for packet inspection.

Preferred embodiments of the invention include in the decision means an expert network much like an expert system but containing an ANN trained with its own implicit rules rather than having explicit rules only. The expert network is adapted to fulfil a number of functions:

1. to interpret clusters of local tentative decisions about FODs into definite decisions on which firm reliable product rejection/classification actions can be taken;
2. to consolidate the information on local tentative decisions from the various input channels feeding it into definite decisions on which firm reliable product rejection/classification actions can be taken;
3. to take prior knowledge into account to help make definite decisions on which firm reliable product rejection/classification actions can be made.

In order to achieve this the expert network preferably includes both ANN components and suitable combinations of more conventional modules such as median, mode and majority decision units, together with other computer learning systems (which might not be classed as ANNs per se). Nevertheless, in a factory installation, it might be important for a line manager

- 6 -

to retain partial or supervisory control of the operation: therefore provision might be incorporated for external adjustment of certain parameters based on additional knowledge or experience. Thus, via suitable input controls the expert network might be temporarily overridden or tuned over a predetermined range for maximum cost-effectiveness, e.g. to make the product rejection rate economic.

Automatic monitoring systems according to the present invention provide the advantage that they are significantly more robust in their decision-making than the individual component analytical processes used within such systems. This is achieved by combining a number of separate component subsystems which work using significantly different methods and then assembling all of the information from the subsystems so that a globally correct and well-informed decision can be taken. Embodiments of the present invention may include amongst other techniques, supervised learning, unsupervised learning, thresholding-based, texture-based, histogram-based and principal component analysis-based methods.

Preferred embodiments of the system also incorporate validation of ANNs by training on large numbers of simulated images. This will overcome a) remanent lack of robustness of ANNs; and b) unfair lack of confidence in ANNs by some parties which sometimes results in them being judged unacceptable for factory use. In particular, a) should permit unsupervised ANN architectures to work acceptably for detecting rare or unknown types of FOD.

A particularly preferred embodiment of the invention includes among its image processing paths modules performing image segmentation based upon: analysis of the spectrum of a histogram of grey-levels within areas of the image; analysis of texture

- 7 -

variations within the image, the analysis including a preliminary step of processing the image to better enhance features of texture (this processing involving convolving the image data with at least one local operator, the local convolution operator preferably having coefficients derived by principal component analysis of the image data itself); analysis of the texture variations within the image using ANNs; and thresholding. In this particular preferred embodiment, the decision means includes an expert network as mentioned above.

A feature of the present invention is that it is capable of detecting an especially wide variety of defects and foreign objects, by virtue of its adaptability and its use of trainable algorithms.

Another feature of the present invention is that it is especially adaptable, and in particular is locally adaptable to conditions in parts of the image, e.g. where the substrate has lower density than normal, so that even in those places it is capable of detecting defects and foreign objects. It is important that it need make no a priori assumptions about the types of defect or foreign object. However, detection can also be improved by tuning the system to special types of defect or foreign object that may arise in a particular application.

Yet another feature of the present invention is that it is especially robust, being both reliable in its detection of defects and foreign objects, while at the same time being robust in eliminating excessive numbers of false alarms, e.g. where the substrate has especially high or low density.

A still further feature of the present invention is that it has especially high sensitivity in the detection of defects and foreign objects. This applies especially in relation to the detection of soft

contaminants, which current systems make little attempt to detect, and where they do are highly unsuccessful at it.

It is believed that there are novel features in the various modules and methods which have been developed in arriving at the particularly preferred embodiment of the invention, as well as in the overall automatic monitoring system and method described above.

Further features and advantages of the present invention will become clear from the following description of the particularly preferred embodiment thereof, given by way of example, and illustrated in the accompanying drawings, in which:

Fig. 1 is a block diagram showing the components of the particularly preferred embodiment of the invention; and

Fig. 2 is a block diagram indicating the structure of a parallel ANN array for use in implementing module 34 of Fig. 1.

The overall FOD detection scheme as illustrated in figure 1 uses a complex hierarchy of different image processing stages which embodies both conventional and ANN techniques. In the following discussion of this embodiment of the invention the system will be described as optimised for the analysis of X-ray images of bags of frozen food (typically bags of sweetcorn kernels) as such bags appear at the end of a factory production line. The FOs which can occur in food products vary widely in size and form ranging from wood, rubber and plastics material to glass, metal and stone. The variation in FO type gives rise to a variation in the X-ray images. Wood, rubber and plastics material (soft contaminants) generally appear lighter than the background food substrate whereas glass, metal and stone (hard contaminants) generally appear darker. No single image analysis technique is powerful enough to detect reliably all possible FOs.

The system is split into four stages: image

- 9 -

capture (1), image pre-processing (2), feature extraction (3) and a decision stage (4).

1. Image Capture Stage

For the purposes of this discussion it will be assumed that the image capture stage involves the taking of an X-ray image with known X-ray apparatus used in automatic monitoring. However, it is to be understood that the images processed by systems according to the invention may be generated by any technology (e.g. infra-red, optical, ultrasonic, etc.) appropriate to the particular application involved.

2. Pre-processing Stage

The requirement of the image pre-processing stage (2) is to reduce artefacts attributable to the imaging technology whilst emphasising the relevant information contained in an image. Dependent on the source of the image, various pre-processing stages will be required (for example averaging (22), log transformation (21) or minimizing), thus increasing the detectability of the various FOs. Where X-ray images are involved log transformation is particularly appropriate because it removes artefacts caused by the X-ray apparatus and compresses the range of grey-levels in the image (which has been found to improve performance).

3. Feature Extraction Stage:

The feature extraction stage is critical and employs a number of different approaches in parallel.

The image processing modules used in the feature extraction stage generally work by characterising the food substrate background and then seeking variations from this background via grey-level thresholding, texture analysis, and the use of artificial neural networks (ANNs). Different types of pre-processing are appropriate to the different types of feature extraction technique used in the feature

- 10 -

extraction stage. Fig. 1 shows connections which make available to each of the modules of the feature extraction stage (3) outputs from each of the pre-processing modules (21, 22 and 23). However, it is to be understood that in a specific application of the system the interconnections between the pre-processing stage and the feature extraction stage can be tailored to the needs of the feature extraction modules.

Various of the individual preferred feature extraction techniques will now be described in greater detail.

3.1 Thresholding methods

Module 31 implements one or more entropy based histogram thresholding methods in order to successfully segment the FOs from the background. At present it is preferred that the module 31 should use a modified entropic measure which decides whether a single or dual threshold is appropriate, dependent on the histogram data.

Thresholding is a popular technique which can be effectively applied to many different image types. It is particularly useful in cases where it simplifies the image content to such an extent that a decision can be made without further processing. However, one drawback to the technique is that there is no single method which can be universally applied to all image types and be expected to produce good results. Instead it is necessary to determine which algorithm is appropriate for the data in question. This decision may initially involve answering simple questions about the histogram data, such as:

Can the histogram be modelled in parametric form or does it require a more complex description?

Will a global threshold suffice or would a series of adaptive thresholds be more appropriate?

Will a single threshold adequately segment the

- 11 -

image or are multiple thresholds required, and if so, how many?

However having answered these questions there are still a range of algorithms which could plausibly be applied, thus it is often necessary to use and compare a number of these methods before deciding on the best for each particular case.

Due to the wide range of potential FO types, and hence the various forms in which they may appear in the image histogram, it was decided to use non-parametric thresholding methods in module 31. These methods are considered to be more robust and more efficient than parametric methods. Adaptive thresholding techniques are particularly valid when there is an underlying variation in the mean grey level value across an image, as happens when uneven lighting conditions occur. For the X-ray images used in developing the present embodiment there was no significant grey level variation across an image, hence adaptive thresholding techniques were not applied. In view of the possible presence of three or more distinct areas in each image (conveyor belt, food substrate, and potentially one or more FO's) both single and dual threshold schemes were applied to the data. It was not necessary to threshold at more than two thresholds since if more than two FO's were present in a bag only one would need to be detected for the bag to be rejected.

Two groups of thresholding techniques were chosen since each uses a different measure of the histogram data to determine the appropriate thresholds. The first method was originally proposed by Otsu [IEEE Trans. Syst. Man Cybern. 9, 62-66 (1979)]. The method is equivalent to minimizing the mean square error between the original grey level picture and its binary representation for a given threshold. The extension of the method to multiple thresholds is simple, although

- 12 -

the thresholds become less credible as the number of classes increases. Variations of Otsu's original method have been proposed by Brink [Pattern Recognition Lett. 9, 355-341 (1989)] and Reddi et al [IEEE Trans Syst. Man Cybern. 14, 661-665 (1984)].

In his original paper, Otsu described three possible discriminant criteria based on ratios of the within-class, between-class and total variance, all of which are equivalent and thus in a given situation any could be chosen. Most researches choose to maximise the between class variance since it is the simplest to calculate.

Using the notation of Otsu an image can be described as being composed of L grey levels. The number of pixels with a given grey level i is termed n_i , and the total of pixels in the image is given by $N = n_1 + n_2 + \dots + n_L$. The probability of a pixel having a given grey level is thus:

$$p_i = n_i / N \quad (1)$$

where

$$p_i \geq 0, \quad \sum_{i=1}^L p_i = 1. \quad (2)$$

- 13 -

For a single threshold the criterion to be maximised is the ratio of the between-class variance to the total variance:

$$= \sigma_B^2 / \sigma_T^2 \quad (3)$$

where the between-class variance is given as:

$$\sigma_B^2 = \omega_0 (\mu_0 - \mu_T)^2 + \omega_1 (\mu_1 - \mu_T)^2 \quad (4)$$

which can be simplified to

$$\sigma_B^2 = \omega_0 \omega_1 (\mu_1 - \mu_0)^2, \quad (5)$$

and the total variance σ_T^2 as:

$$\sigma_T^2 = \sum_{i=1}^L (i - \mu_T)^2 p_i \quad (6)$$

Since the total variance (equation 6) is constant for a given image histogram, equation 3 simplifies to maximising the between-class variance. This can be done using the following definitions:

$$\omega_0 = \sum_{i=1}^k p_i, \quad \omega_1 = \sum_{i=k+1}^L p_i = 1 - \omega_0, \quad (7)$$

$$\mu_0 = \sum_{i=1}^k i p_i / \omega_0, \quad \mu_1 = \sum_{i=k+1}^L i p_i / \omega_1. \quad (8)$$

The method can easily be extended to the dual threshold case $1 \leq k_1 \leq k_2 \leq L$, where the resultant classes are given by C_0 , C_1 and C_2 , with grey level ranges of $[1, \dots, k_1]$, $[k_1 + 1, \dots, k_2]$ and $[k_2 + 1, \dots, L]$ respectively. The dual thresholds are found by maximising one of Otsu's three original criteria as before.

In the development of module 31 both the single threshold scheme given by equation 3 (hereafter referred to as method 1), and the dual threshold scheme (referred to as method 2) were used.

There are a number of problems with Otsu's method. Of particular importance in the present context is the effect whereby certain object sizes, and hence

- 14 -

population ratios in the histogram, can lead to the selection of an incorrect threshold. In the context of FOD detection this would mean that certain sizes of FO's may be too small to be detected.

The second group of thresholding algorithms use the concept of the entropy of a histogram to determine the appropriate threshold. The algorithms used in development of module 31 are based on those of Kapur et al [Comput. Vision Graphics Image Process. 29, 273-285 (1985)] which, for a single threshold scheme separate the histogram grey level probabilities into two distributions, one associated with the foreground and the other with the background of the image. The entropy of each of these distributions is then combined and the grey level with the maximum total entropy is taken as the appropriate thresholding position.

For the single threshold case Kapur et al. divided the probability distribution into two classes; those with grey levels up to the threshold value k , $[1, \dots, k]$, and those with grey levels above $[k+1, \dots, L]$. This led to two probability distributions given by:

$$\text{Class A: } \frac{p_1}{P_k}, \frac{p_2}{P_k}, \dots, \frac{p_k}{P_k} \quad (9)$$

$$\text{Class B: } \frac{p_{k+1}}{1-P_k}, \frac{p_{k+2}}{1-P_k}, \dots, \frac{p_L}{1-P_k} \quad (10)$$

where

$$P_k = \sum_{i=1}^k p_i \quad 1-P_k = \sum_{i=k+1}^L p_i \quad (11)$$

The entropies for each class are given by:

$$H(A) = - \sum_{i=1}^k \frac{p_i}{P_k} \ln \frac{p_i}{P_k} \quad (12)$$

- 15 -

$$H(B) = - \sum_{i=k+1}^L \frac{P_i}{1-P_k} \ln \frac{P_i}{1-P_k} \quad (13)$$

And the total entropy is defined as

$$\Psi(k) = H(A) + H(B). \quad (14)$$

Expansion of which leads to:

$$\Psi(k) = \ln \left(\sum_{i=1}^k p_i \right) + \ln \left(\sum_{i=k+1}^L p_i \right) - \frac{\sum_{i=1}^k p_i \ln p_i}{\sum_{i=1}^k p_i} - \frac{\sum_{i=k+1}^L p_i \ln p_i}{\sum_{i=k+1}^L p_i} \quad (15)$$

Equation 15 is hereafter referred to as method 3.

The extension of the technique to dual thresholds is given by:

$$\Psi(k_1, k_2) = \ln \left(\sum_{i=1}^{k_1} p_i \right) + \ln \left(\sum_{i=k_1+1}^{k_2} p_i \right) + \ln \left(\sum_{i=k_2+1}^L p_i \right) - \frac{\sum_{i=1}^{k_1} p_i \ln p_i}{\sum_{i=1}^{k_1} p_i} - \frac{\sum_{i=k_1+1}^{k_2} p_i \ln p_i}{\sum_{i=k_1+1}^{k_2} p_i} - \frac{\sum_{i=k_2+1}^L p_i \ln p_i}{\sum_{i=k_2+1}^L p_i} \quad (16)$$

and is referred to as method 4 in the following.

In Kapur's method the entropy of the histogram is based on Shannon's definition of entropy which may give rise to a number of problems. One of which is that Shannon's entropic description is not defined for distributions which include probabilities of 0. For this reason amongst others, a number of researches have proposed variations on Kapur's original method, including: Pal et al [IEE Proc. Pt.E 136, 284-295 (1989) and IEEE Trans. Syst. Man Cybern. 21,1260-1270 (1991)]; Abutaleb [Comput. Vision Graphics Image Process. 47, 22-32 (1989); and Brink [Pattern Recognition 25, 803-808 (1992)].

Several authors have undertaken comparisons of the various thresholding methods, such as; Lee et al [Comput. Vision Graphics Image Process. 52, 171-190 (1990)]; Sahoo et al [Comput. Vision Graphics Image

- 16 -

Process. 41, 233-260 (1988]; Tsai et al [pattern recognition Lett. 13, 245-252 (1992)]; and Abutaleb himself. Sahoo et al concluded that both Otsu's and Kapur's methods worked well, although Otsu's was slightly better. Abutaleb subjectively compared the performance of his two-dimensional entropic threshold algorithm with that of Kapur et al and showed a reduction in the amount of noise present in the resultant images when the average grey level value was included. For many cases this is desirable. However, in the context of detection of FOs it is only required that any FO's affecting a substrate should be detected, it is not required that the location of the boundary between the object and the substrate should be accurately located. Thus the sensitivity of the method is more important in the present context.

During the development of the present embodiment the effectiveness of Otsu's and Kapur et al's thresholding methods have been evaluated for both single and dual threshold cases. A modified entropy method has been developed which gives an improved performance compared with the previous techniques. In preferred embodiments of the invention it is this modified entropy method which is implemented by the thresholding module 31.

In the context of FO detection is bagged food it has been found worthwhile to apply both methods 3 and 4 mentioned above. However, by further processing the entropy measure of method 3, the alternative method was devised. The new method has an advantage over methods 1 to 4 in that it proposes single or dual thresholds dependent on the information characterised by the entropy measure. The significance of this feature is that it gives a clear indication as to the presence or absence of FO's in an image.

It has been found that for an X-ray image of a bag of frozen food plus FO resting on a conveyor the

- 17 -

histogram shows a sharp peak at the high grey level values associated with the background conveyor belt. There is also a broader peak, produced by the frozen food, in the mid to lower grey level range. The peak associated with the small dark FO does not show at this scale.

The corresponding entropy measure shows a major, broad peak associated with the frozen food grey levels. This is the peak detected by method 3. Also there is a second sharper peak at lower grey levels. This peak is associated with the FO. In view of the size of the second peak it can be said to contain important information about the image and represents a plausible alternative level to threshold at.

The new method proposed here searches the entropy function for alternative peaks of this sort, and then thresholds at the single or dual thresholds as appropriate.

A number of possible strategies can be followed when searching for alternative peaks, it is preferred to use the following:

- 1) Determine the major peak (k_1) using method 3.
- 2) Search from the lowest populated grey level up to k_1 for the lowest alternative peak below k_1 (L_T).
- 3) Search from the highest populated grey level down to k_1 for the highest alternative peak above k_1 (H_T).
- 4) If L_T and H_T both exist, determine which is the major (k_2).
- 5) If only one of L_T and H_T exists assign it to k_2 .
- 6) Threshold image at k_1 and k_2 as appropriate.

This has the following desirable properties:

- i). It does not force an image to be thresholded at two values unless two thresholds are appropriate.
- ii). It is faster to compute than method 4.

- 18 -

iii). It can be made sensitive to small FO's which are often missed by alternative thresholding methods.

Variance based methods proposed by other authors yield good results which are probably due to the relatively large size objects found in their images (since it has been found that methods 1 and 2 are not sensitive enough to detect the small FO's). It has also been found that entropic thresholding is a more appropriate method for detecting objects up to at least a box size of 10.0 % of the total image size. This is in accordance with the results of other researchers and highlights the need to choose a thresholding method appropriate to the size of the objects expected in the images.

The new method proposed here (method 5) showed several improvements over the other entropy based measures (methods 3 and 4). The major, one, from the point of view of detecting FO's, is that it does not produce dual thresholds unless they are relevant to the histogram data. This removed the need for further analysis to determine whether the threshold is due to the presence of a FO or not. The sensitivity of the method can be easily altered by changing the secondary peak detection scheme to recognise smaller peaks in the entropy measure, although, obviously, as the sensitivity is increased so will the rate of false object detection. It should be noted that for the food industry it is better to be over cautious when setting the sensitivity of a FO detection method than to let through a large number of faulty goods. It may be possible to further improve the results of method 5 by forming the histogram from a subset of lines in the image being examined.

Method 5 proposed here has been found to give results which are similar to those obtained by human thresholding. This suggests that they are close to the best obtainable via thresholding methods. Since this

- 19 -

still does not give 100 % object detection it is necessary to use alternative strategies (i.e. implemented by the other feature extraction modules) to increase the rate of FO detection.

3.2 Texture methods

In modules 32 and 33 of the preferred embodiment shown in Fig.1, the emphasis is on accentuating features of texture and detecting, via module 38, any differences highlighted from the texture analysis.

Various methods are known for characterising texture in images (i.e. processing image data so as to emphasise features of texture and differentiate between areas of an image having different textures). The known methods include convolving image data with local operators, Fourier methods and use of co-occurrence matrices.

Several researchers have used convolution filter methods for texture characterisation. The convolution masks used in these methods are designed to act as matched filters for certain types of variations found in textures.

Laws (1980) defined an empirical set of masks and used them to generate feature images. These masks were derived from one dimensional vectors which are sensitive to characteristics such as edges and spots. Ade (1983), presented a method where mask coefficients are adaptively generated via the calculation of eigenvectors. Unser (1986) proposed a filter bank analysis approach, where the problem of texture analysis was approached by taking some local linear transform such as the discrete cosine transform over each sub-image.

More recently several researchers (Turner, 1986; Fogen and Sagi, 1989; Jain and Farrokhnia, 1992) have used Gabor filters as convolution masks. The motivation there was to show that the receptive field

profiles of certain simple cells in the visual cortex approximate the Gabor functions. However these methods have not proved to be any better with respect of recognition or computational needs than existing texture analysis methods (Ohnian and Dubes, 1992).

In the preferred embodiment of the invention illustrated by Fig.1, one texture module (32) uses convolution masks sensitive to the particular local structural attributes, and another texture module (33) uses Laws masks, in order to highlight texture features. Based upon these features, the image is segmented by module 38 into regions corresponding to the substrate and regions corresponding to an FOD.

As mentioned above, module 32 uses convolution masks sensitive to the particular local structural attributes. For texture characterisation, it is preferred to use masks having coefficients which have been "learnt" from an example of the food substrate with no contaminants. These filters respond to the peculiarities of the structure of the texture and subsequently detect any anomalies, such as FOs within the food substrate. It is presently preferred to use principal component analysis (PCA) techniques to find orthogonal vectors in data space that account for as much as possible of the variance of the image data. The vectors or principal components are then used as the coefficients of the convolution masks.

Two methods may be used for the generation of principal components; one based on the calculation of the variance-covariance matrix and the other using an ANN model.

Below there is a brief mention of the various texture analysis methods using PCA that have been proposed and a more detailed description of an artificial neural network (ANN) technique which it is preferred to use in module 32 in order to perform the PCA and obtain one or more local convolution operators.

After calculating the variance-covariance

- 21 -

matrix of a homogeneous texture region Ade (1983) obtained the mutually orthogonal eigenvectors of the matrix. These vectors best characterise the structure of the underlying texture and are used as coefficients of the masks. His approach maybe criticised for its computational demand. This is particularly true for large input as this requires eigenvector calculation of a higher dimension.

A 2-layer (1 hidden layer) Multi-Layer perceptron (MLP) using the Back-Propagation (BP) rule can be trained to perform an identity mapping with fewer nodes in the hidden layer than input nodes; this network can be used to approximate the principal components (Cotterell, Munro and Zipser, 1987). The hidden units project onto the subspace of the M principal components. The major difference as indicated by Sanger (1989) is that an ANN trained using the Generalised Hebbian Algorithm (GHA) produces a linear combination of the M eigenvectors. Other factors regarding the BP algorithm such as: training times, local minima, selection of the number of hidden nodes, further reduces the usefulness of the MLP.

The GHA algorithm mentioned above was developed to train networks in order to find the eigenvectors of the autocorrelation matrix of input data, given only samples from the input distribution. Each output of the trained network represents the response to one eigenvector and outputs are ordered in decreasing value. It is preferred that module 32 should include an ANN trained using the GHA algorithm in order to perform PCA and produce the convolution masks for use in texture characterisation. The ANN finds the eigenvectors of the correlation matrix, but the principal components are the eigenvectors of the covariance matrix. For zero mean data there is no difference. The method does not require the calculation of the covariance matrix, since the eigenvectors are

- 22 -

derived directly from the data. In some cases actual principal components generated by Ade's methods are more suitable than the approximations generated more rapidly by ANN. This has a considerable computational speed and memory advantage over conventional methods particularly if the input dimensions are large. The GHA is given by:

$$W_{nq}(t+1) = W_{nq}(t) + \eta V_q \left(X_n - \sum_{j=1}^q V_j W_{nj} \right)$$

$$V_q = \sum_{n=1}^N W_{nq} X_n$$

where V_q is the output at each node and W_{nq} are the inter-connection weights. The learning rate η can decrease with time or can be held constant. The outputs project the input vector X , of dimension N , onto the space of the first M principal components.

In developing module 32 an ANN topology of the above type was having 9 input nodes and 8 output nodes. This required training. The set of training data that was used contained 6600 samples, extracted from a 3 x 3 mapping window which scanned across a small section of an image with no FO. The learning rate η was held constant at 0.1, although it could be made to decrease to zero as t increases. The weights which the network learned are represented as eight, 3 x 3 masks. The masks are the rows W_n of the 8 x 9 weight matrix. Larger mask sizes of 5 x 5 to 15 x 15 have also been used in order to maximally reflect the structure of the underlying texture.

The training data for the GHA contained 6600 samples and was iterated once through the network with the learning rate held constant. The ANN generalisation achieved using these parameters was adequate.

- 23 -

achieved using these parameters was adequate.

Once the ANN has generated coefficients as described above, the images are then convolved with masks whose coefficients are the principal components generated using the ANN. The convolution transforms the texture segmentation problem from the spatial domain to the feature space domain.

The masks generated by the ANN are not all equally sensitive to the textural variations in the images. Thus it is preferable to generate a set of possible masks, test their performance during a set-up stage using typical test data and select a sub-set of the masks for use. The image data can then be convolved with each of the sub-set of masks to provide a number of images with enhanced texture features. In general the output from the higher energy mask tends to hide the discriminatory information between the FO and the background food substrate. A method to overcome this while maintaining sensitivity (performed by module 39) is to reduce the higher energy coefficients, e.g. in proportion to the corresponding Eigenvalues, according to an appropriate theoretical model, or by some method based on experiment or practical experience or knowledge, or by the output of a separate ANN learning module: such methods might be wholly or partly automatic, e.g. partially under the control of a factory line manager, for optimal cost-effectiveness.

When testing the above-described ANN using X-ray images of bags of frozen corn kernels it was found that masks 5, 7 and 8 generated the most discriminable features. Masks 5 and 7 could be seen as diagonal edge detectors in opposite orientation, and mask 8 could be seen as a ripple detector. The three masks were then used in parallel to generate feature images from input image data. These feature images were then smoothed to

- 24 -

remove any noise. The resulting feature images can then be further processed to enable detection of any variation on the textured food substrate as a different output from the general learnt texture area. Thresholding can then be used to segment the FO from the food. However, in the post-processing stage before the thresholding, it is advantageous to use a maximising filter to enhance the contrast between any detected FO and the filtered food substrate. Finally, after segmentation, the three resulting images were then passed through a majority operator where the final pixel assignment corresponded to the majority label of the three pixels in the segmented images.

It is advantageous for the above-described texture module 32 to receive image data which has passed through a log filter in the pre-processing stage. This increases the grey-level contrast between the FO and the food substrate. This is particularly useful for hard contaminants where the FOs appear darker than the food substrate.

The results obtained using this embodiment of texture module 32 have been promising. This type of texture characterisation is based on a convolution operator which "adapts" and "learns" from representation of food samples containing no contaminants using the generalised Hebbian rule. The detected FOs are discrepancies within the food substrate which are highlighted by the convolution process. Such a method offers a fast and computationally efficient method for calculating principal components of the input distribution.

Module 33 is similar to module 32 but, in this embodiment, uses Laws masks as the local convolution operator. During set-up of the apparatus a selection is made of a sub-set of Laws masks which most clearly

- 25 -

characterise the texture of the particular substrate that is to be used. The image data is then convolved with each of the selected Laws masks and the resultant feature images are processed to arrive at a single segmented image as in the case of module 32.

Texture modules 32 and 33 generate feature images which may then be segmented (i.e. processed to delineate homogeneous areas of different types) in any of a number of ways. Fig. 1 illustrates that the segmentation (performed by module 38) could be performed by an unsupervised ANN or a thresholding technique.

3.3. Artificial Neural Networks

Research into artificial neural networks (ANNs) is undertaken in the hope of achieving human-like performances for pattern recognition tasks. The pattern recognition application addressed here is the detection of foreign contaminants in food bags. Several ANN architectures exist [Simpson, 1990], each characterised by the particular connection topology of their nodes (processing elements) and learning rule. Learning is a task whereby the connection weights between the nodes are adapted such that the network can undertake a specific task.

There are two categories of learning: supervised and unsupervised. In supervised learning, there is a supervisor to teach the system how to classify a known set of training patterns, i.e. both the training patterns and the associated desired output patterns are available. In unsupervised learning on the other hand, all one has is a collection of patterns and the network usually categorises them according to

- 26 -

some prespecified measure of similarity between them. The unsupervised approach is attractive because the networks learns without assuming much knowledge about the data. However, ANN 34 using supervised learning will first be described.

As mentioned above, the analytical method employed by texture module 32 is particularly good at detecting hard contaminants such as stone, metal and glass. However soft contaminants have similar X-ray attenuation to the organic foods and consequently the grey level contrast in the image is low. This makes the detection of soft contaminants more difficult.

Module 34 uses a supervised paradigm where, during an initial set-up phase, a supervisor teaches the network how to classify a known training set of samples. After several iterations of the training data the network develops its own features which represent both the FOs and the food substrate.

Representing a wide variety of FOs to an ANN to build an adequate generalisation set is a complex task. As the number of FOs that have to be learnt increases, training complexities and times increase and furthermore there is a possibility that the ANN might not converge to a minimum. When seeking to detect soft contaminants it has been found to be advantageous to create sub-networks, each one being trained to recognize one particular FO. Using this parallel ANN (PANN), learning is improved and thus a better representation of each FO is generated. It has been found that ANNs comprising a few nodes can successfully detect FOs.

In developing supervised ANN module 34 each sub-network was implemented using a fully connected multi-layer perceptron (MLP) topology with the Back-Propagation (BP) learning rule [Rumelhart and McClelland, 1986], which comprised a three layer topology with one layer of hidden nodes. It has been

- 27 -

is possible to train MLP networks with one hidden layer using the BP training and sigmoidal non-linearities as output functions to form complex regions. The network is given the input and the corresponding desired output pairs and it learns to minimise the error between the network output and the desired output by automatically adjusting the inter-nodal connection weights.

In this example, the input layer nodes have a linear transfer function whereas the input-output characteristics of all other nodes are defined by a sigmoidal non-linearity threshold function:

$$f(x) = \frac{1}{1 + e^{-x}}$$

Each layer in the MLP topology is a column of nodes with feed-forward communication between layers. Each layer calculates:

$$y_j = f \left[\sum_{i=1}^n w_i x_i \right]$$

and passes that as input to the next layer. The final layer output values are o_j and the desired output values are d_j .

The BP algorithm minimises the difference between the desired output and the actual output values:

$$E = \sum_p \sum_j (o_{pj} - d_{pj})^2,$$

for all nodes in the output layer; $j \in (1, 2, \dots, J)$ and the input patterns; $p \in (1, 2, \dots, P)$. The function is minimised by adjusting the inter-nodal connection weights and the nodal thresholds. The detailed algebra BP rule is not presented here, but can be found in Rumelhart and McClelland [1986].

Each sub-network from the overall PANN architecture is trained to recognise a specific FO. A typical sub-network may consist of an input layer of

- 28 -

network is trained to recognise three regions on the images: foreign object, food substrate and conveyor belt background. The input layers of nodes each yield values from a 3 x 3 mapping window which scans across the image. The choice of hidden layer nodes for each subnet is problematic; if there are too many the network will memorise the training samples and consequently not generalise. On the other hand, if there are too few, then the network will not converge during training. It is preferred to use a dynamic node creation method which automatically chooses the optimum number of nodes in the hidden layer.

Classification accuracy for segmentation purposes implies a comparison of the segmented region with that of a visually perceived segmented image. However, it is not possible to draw out a visual boundary segmenting the FOs from each corrupted image. Consequently, data is extracted via a 3 x 3 mapping window which scans across small subsections of the three categories of a training image. A cross-validation procedure is used where the data is split into training data and test data, and the latter is used to measure the generalisation of the network. The nodes in the middle layer are varied between a lower and an upper bound using the dynamic node generation method in order to achieve good segmentation. During training, the node outputs are set to 0, except for the class from which the current input pattern is taken, which is set to 1. The patterns are iterated through the network until sufficient convergence is achieved. After training, the internodal connection weights and nodal offsets are frozen and the network used to detect FOs. The node with the highest response is the label assigned to the pixels being classified.

It is advantageous for the image data input to supervised ANN module 34 to undergo averaging in the

- 29 -

cluster the grey levels; this tends to connect the neighbouring pixels.

In a trial, three networks each representing a category of FO were set up in the PANN architecture of figure 2. Training and test data was generated using X-ray images of FOs in bags of frozen corn kernels on a conveyor. In the trial, one of the networks was trained to recognise a rubber grommet, another to recognise wood and another to recognise a rubber eraser. Based on the node generation method, the best generalisation achieved for each of the FO sub-nets were 9-6-3, where the number of hidden layer nodes was 6. Good classification rates were achieved by the three sub-nets (97.3 % for wood, 92.87 % for the grommet). The individual sub-net segmented images were filtered through a noise removing filter. The time taken to segment 256 x 256 x 8 resolution images was approximately 45 sec. - unaided by any special hardware. All simulations were implemented on a DEC 5000/25 workstation. (For all practical purposes in the factory environment, it is envisaged that the networks would be trained over-night.)

The wood piece used in the trial was approximately of dimension 48 x 18 x 17 mm: smaller pieces could not be substantially discriminated from the food substrate. Furthermore, the edge of the bag where the layers of corn kernels decreased tended to be mistakenly classified as the FO. This is because the X-rays absorbed by the corn is similar to that absorbed by the FO. However, the threshold decision unit can cope with such an overall discrepancy. For plastics, we have successfully detected the rubber grommets and erasers. However, the detection of the eraser is straightforward and can be done using thresholding techniques. Again, plastics or polymers of low density such as a piece of vehicle indicator lamp reflector, cannot be seen on the X-ray image even with a trained human eye and has thus

- 30 -

to enhance the contrast between the contaminants and the food are being suggested. The possibility of doping likely contaminants with a high atomic number marker is likely to increase to increase detectability to a useful degree.

The ANN of module 35 uses unsupervised learning, all one has is a collection of samples and the network categorizes them according to some measure of similarity between them.

In the above discussion of texture methods it has been described how segmentation of images can be based on textural features and a logical ANN paradigm using an unsupervised learning rule. Visa [1992], also segments images using texture features in an unsupervised mode using a Kohonen self-organising topological feature map. The unsupervised learning rule in our ANN is similar to Kohonen's vector quantisation methodology [Kohonen, 1990] and briefly is as follows:

step 1. Select a pattern $X_p = (x_1, x_2, \dots, x_N)$ from the training set; $p \in (1, 2, \dots, P)$ and feed into input layer. N is the dimension of the pattern vector.

step 2. Find node W_j ; $j = (1, \dots, J)$ closest to X_p using:

$$\|X_p - W_1\| = \min_{j=1}^J \|X_p - W_j\|$$

where W_j is the closest node to X_p and the measure of similarity is the Euclidean distance.

- 31 -

step 3. Move W_1 closer to X_p using:

$$W_1(t+1) = W_1(t) + \alpha(t)[X_p(t) - W_1(t)] ,$$

where $\alpha(t)$ is the learning rate at time t .

step 4. Repeat steps 1 to 3 for successive instants of time t until convergence.

The ANN consists of a topology of nodes which represent both the area on the bag which does not contain any FOs and the FOs. Although this method analyses textured images, it is not dependent on a mathematical measure for texture and therefore eliminates the long processing times which might be needed to characterise the texture. For this type of feature extraction process it is advantageous for the pre-processing to use an averaging filter to remove spot noise and cluster the grey levels; this tends to connect neighbouring pixels.

The training data for the node which represents the part of the bag which is free from contaminants may be acquired from a 3 x 3 mapping window and is then represented by a 9-element vector. The variation within such images over a small 3 x 3 area can be substantial because the X-rays may penetrate an uneven layer of food. In order to accommodate such a variation, more than one node is used to represent the food substrate. The vector quantisation learning procedure is used to adapt the data, until the nodes sufficiently represent the food substrate.

After the nodes representing the food substrate have learnt the essential features thereof, nodes representing the FOs need to be characterised. There exists a large number of different kinds of impurities that might be found in the food bags and therefore no a priori knowledge of the FOs can be used.

- 32 -

contaminants appear at either the darker or lighter end of the grey level spectrum (this difference is not always distinct enough for detection using simple thresholding techniques). We therefore have two nodes at each end of the spectrum representing the hard and soft contaminants.

After the nodes have been sufficiently adapted, the test images are scanned by a 3 x 3 window and the 9-element data vectors are passed through the network. The vectors are assigned the label of the node which best matches the input. The Euclidean distance is used as a measure of similarity between the pattern sample under test and the nodes. The clusters of regions formed around the nodes representing the food substrate can be merged into one homogeneous region. Here unsupervised networks can reveal regions on an image which might otherwise have been unobserved using supervised networks. These clusters can be of considerable significance as will be explained below. The method has been tested on several images and works well. For example, this technique enables the detection of a small rubber grommet embedded in a bag of frozen corn kernels.

Techniques for the detection of a specific FO, for example wood, via texture recognition and a Multi-Layer Perceptron topology using the supervised Back Propagation learning rule are discussed above in connection with the texture modules.

3.4 Histogram spectral analysis

In module 36 a histogram is produced using the grey levels contained in a number of lines of the (possibly pre-processed) image. This histogram can be thought of as a spectrum of the grey levels. Any FOs show up as a change in the spectrum. This spectrum can be analysed using Fourier techniques, ANN's or other analysis techniques to detect these changes.

3.5 Boundary shape analysis

The FOs when embedded near the edge of the food are extremely difficult to detect using the above mentioned techniques. Module 37 uses analysis of a series of food substrate boundaries to improve the detection of FOs located close to the edge of the food.

Decision stage

The various techniques used in the feature extraction stage result in a number of "feature images" which provide information about particular features of the original image. Based upon this information a decision is to be made as to whether the food sample contains a FO or not. This can be done by several methods such as: an ANN which embodies implicit rather than explicit rules, and which could be trained on extra inputs from conventional modules such as a majority decision device. Alternative decision modules which could be used in place of or in addition to an ANN are rule-based systems.

Although a specific embodiment of the invention has been described it is to be understood that many features of the specific embodiment may be varied without departing from the present invention. For example, the number and type of image processing modules used in the feature extraction stage may be altered. One particular modification in that respect which is believed to be advantageous is to include in the feature extraction stage a texture module using PCA implemented by an ANN and a texture module using PCA implemented by calculating the variance-covariance matrix.

CLAIMS:

1. An automatic monitoring system for detecting foreign objects or defects affecting a substrate, the system comprising:

imaging means for obtaining an image of at least a portion of the substrate;

image processing means comprising a plurality of parallel image processing paths applying different analytical techniques to evaluate which regions of the image correspond to foreign objects or defects; and

decision means receiving inputs from the parallel image processing paths and adapted to decide whether or not the imaged portion of the substrate is affected by a foreign object or defect.

2. An automatic monitoring system as claimed in claim 1, employing trainable means comprising supervised or unsupervised artificial neural networks.

3. An automatic monitoring system as claimed in claim 1 or 2, employing a plurality of cognate statistical pattern recognition methods including threshold based, texture based, histogram based, boundary analysis based and principal component analysis-based techniques.

4. An automatic monitoring system as claimed in claims 1, 2 or 3 in which the decision means comprises an expert neural network containing an artificial neural network trained with its own implicit rules, interpreting clusters of local tentative decisions about foreign objects and defects into definite decisions.

5. An automatic monitoring system as claimed in claims 1, 2, 3 or 4, in which one or more of the processing paths comprises an artificial neural network.

- 35 -

6. An automatic monitoring system as claimed in any of claims 1-5, comprising among its image processing paths, modules performing image segmentation based upon analysis of texture and intensity histogram variations within the image, the analysis including a preliminary step of processing the image to better enhance features of texture.

7. An automatic monitoring system as claimed in any of claims 1-6 comprising an image pre-processing stage, reducing artefacts attributable to technology employed in the imaging means whilst emphasising the relevant information contained in an image.

8. An automatic monitoring system as claimed in any of claims 1-7, arranged so as to identify foreign objects in prepacked containers of food.

9. An automatic monitoring system as claimed in any of claims 1-8 comprising imaging means capable of processing images appropriate to the particular application e.g. infra red, optical, ultrasonic etc.

10. A method for automatically monitoring a substrate for foreign objects or defects affecting the substrate, the method comprising the steps of;
imaging at least a portion of the substrate;
processing the captured image using a plurality of parallel processing paths applying different analytical techniques to evaluate which regions of the image correspond to the normal substrate and which regions of the image correspond to foreign objects or defects; and

making a judgement as to whether or not the imaged portion of the substrate is affected by a foreign object or defect, the judging step being influenced by

- 36 -

outputs from the parallel processing paths.

11. The method according to claim 9, employing supervised or unsupervised artificial neural network techniques.

12. The method according to claims 9 and 10, employing a plurality of cognate statistical pattern recognition methods including threshold based, texture based, histogram based, boundary analysis based and principal component analysis-based techniques.

13. The method according to claims 9, 10 or 11 employing an expert neural network trained with its own implicit rules, interpreting clusters of local tentative decisions about foreign objects and defects into definite decisions.

14. The method according to any of claims 9-12 wherein any of the parallel processing paths employs an artificial neural network.

15. The method according to any of claims 9-13 performing among its parallel processing paths, image segmentation based upon analysis of texture and intensity histogram variation within the image, the analysis including a preliminary step of processing the image to better enhance features of texture.

16. The method according to any of claims 9-14, employing an image pre-processing stage, reducing artefacts attributable to the imaging step whilst emphasising the relevant information contained in an image.

- 37 -

17. An automatic monitoring system substantially as hereinbefore described and with reference to the accompanying drawings.




Figure 2




Figure 1

INTERNATIONAL SEARCH REPORT

Int'l Application No

PCT/GB 94/02698

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06T7/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	G.S.DESOLI ET AL 'A system for automated visual inspection of ceramic tiles' November 1993 Proceedings of the IECON '93, international conference on industrial electronics, control and ... , MAUI, HAWAII , XP437519	1,10
A	see page 1875, column 1, line 21 - page 1875, column 2, line 2; figure 9 --- -/--	3

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

3 April 1995

Date of mailing of the international search report

12 04 95

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Anderson, A

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 94/02698

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	MACHINE VISION APPLICATIONS, ARCHITECTURES, AND SYSTEMS INTEGRATION II, PROCEEDINGS, vol.2064, September 1993, BOSTON, USA pages 124 - 134 B.G.BATCHELOR 'Automated inspection of bread and loaves' see page 124, line 46 - page 125, line 27 see page 125, line 43 - page 125, line 45 ---	8
A	MESURES, 17 October 1988, PARIS, FRANCE pages 59 - 61 'Corps etrangers en iaa; quand les rayons x traquent l'indesirable' see page 61, column 2, line 7 - page 61, column 3, line 22 ---	1,8
A	US,A,5 260 871 (GOLDBERG) 9 November 1993 see abstract -----	2

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 94/02698

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A-5260871	09-11-93	NONE	